

GPU-Accelerated Deep Learning for Heatwave Prediction and Urban Heat Risk Assessment

Adis Alihodzic (ORCID: 0000-0003-0761-1667)^{1*}

^{1*}Department of Mathematical and Computer Sciences, Faculty of Science, University of Sarajevo, Sarajevo, Bosnia and Herzegovina.

Corresponding author(s). E-mail(s): adis.alihodzic@pmf.unsa.ba;

Abstract

Heatwaves are an important problem in cities, and climate change makes this problem more difficult. In this paper, we present a GPU-based deep learning framework for next-day prediction of urban thermal conditions and for heat risk assessment. The study was carried out in Sarajevo by using MODIS land surface temperature data and Open-Meteo forecast data. We tested several models, including convolutional models and spatiotemporal models. Among them, ConvLSTM with a mixed loss function gave the best results. The obtained values were MAE = 0.2293, RMSE = 0.3089, and $R^2 = 0.8877$. The experiments also showed that results can be improved by using longer temporal series and additional meteorological variables. Since the framework was implemented on a GPU and trained with mixed precision, the execution time was reduced. Based on the predicted temperature fields, it was also possible to combine hazard information with exposure and vulnerability data in order to generate city heat risk maps. The proposed framework can be used as a practical basis for city heat analysis.

Keywords: heatwave prediction, urban heat risk, deep learning, GPU acceleration, spatiotemporal models, ConvLSTM, climate data analysis

1 Introduction

Heatwaves are the most effective climate extreme today. It is well known that they have become more frequent, more severe, and longer in many regions. Their effects can be seen in public health, energy demand, labor productivity, and urban infrastructure. Recent studies have revealed that heat-related mortality is strongly associated with extreme hot periods, especially in cities, where the urban environment can further increase thermal stress [1–3]. Also, urban areas are particularly sensitive to extreme heat, as dense construction, impervious materials, reduced vegetation, and limited air circulation often make cities warmer

than surrounding rural areas. This phenomenon is known as the urban heat island. During heatwaves, the combined effects of climate change and urban heat island effects may aggravate thermal conditions. Therefore, predicting urban heat conditions and estimating urban heat risk have become essential topics in environmental modeling, public health, and urban planning [1, 2, 4]. Another crucial point is that urban heat does not affect all elements of a city equally. Some regions of the city have more vegetation and better ventilation, while others are represented by dense built-up areas and large impervious surfaces. Population exposure is also not uniform across all locations. Recent studies have shown that urban

heat exposure may be spatially uneven and closely linked to demographic and socioeconomic patterns. Because of this, urban heat analysis should not be based solely on temperature predictions but should also incorporate exposure and vulnerability factors [3–6]. Satellite products provide spatially detailed land surface temperature information, while gridded population datasets can be used to estimate population exposure [7, 8]. MODIS land surface temperature products are widely used in thermal remote sensing and have been extensively validated in the literature [9, 10]. Population data from the WorldPop project are also often used in spatial exposure analysis [8, 11]. Conventional heatwave analysis has commonly relied on statistical methods, physical modeling, and conventional machine learning approaches [12, 13]. Although these methods are appropriate, they may be limited when the process exhibits strong spatial and temporal dependencies simultaneously. For that reason, deep learning methods have received increasing attention in recent years because they have shown good results in spatiotemporal prediction problems [14, 15]. Convolutional neural networks are suitable for learning spatial patterns [16], while ConvLSTM architectures are designed to model spatial and temporal structure jointly. Besides prediction accuracy, computational efficiency is also important in practice. High-resolution urban heat modeling involves large spatiotemporal datasets and repeated model training, which may require substantial computational time [17, 18]. Therefore, GPU acceleration plays an important practical role, enabling faster training and inference in modern deep learning workflows [19, 20]. This is especially useful when the forecast results need to be transformed into urban heat risk layers for practical decision support [5, 21]. In this chapter, heatwave prediction and urban heat risk assessment are considered from a practical perspective. The main idea is to combine satellite-derived land-surface temperatures with daily meteorological forcing within a GPU-accelerated deep learning framework. Two model families are considered, namely a CNN baseline and a ConvLSTM model. In addition, a simplified urban heat risk layer is constructed by combining predicted thermal intensity with exposure and vulnerability information. In that way, this chapter connects temperature prediction with a risk-based interpretation that can

be useful in climate adaptation and urban planning studies [5, 6]. The main contributions of this chapter are as follows. First, a practical framework for combining satellite-derived thermal data and daily meteorological forcing is presented for short-term urban heat prediction. Second, the role of GPU acceleration in training and evaluation of deep learning models is discussed. Third, the influence of dataset design, temporal coverage, and multi-location meteorological forcing on predictive performance is experimentally analyzed. Fourth, a simplified methodology for generating urban heat risk maps from predicted thermal conditions, exposure factors, and vulnerability factors is outlined. The remainder of this chapter is organized as follows. Section 2 presents related work on heatwave prediction, urban heat analysis, and deep learning approaches. Section 3 describes the study area, data sources, and preprocessing steps. Section 4 introduces the proposed GPU-accelerated framework and the prediction models considered. Section 5 explains the construction of the urban heat risk layer. Section 6 presents the experimental results and performance analysis. Finally, Section 7 gives the concluding remarks and possible directions for future work.

2 Related Work

2.1 Deep learning for heatwave and urban thermal prediction


Recent studies have shown growing interest in deep learning for heat-related prediction tasks. In the broader area of extreme heat forecasting, deep learning models have already been applied to predict extreme heat events from meteorological variables, and the reported results indicate that such data-driven approaches can achieve good predictive performance [22]. Deep learning has also been used in urban thermal studies, especially for land surface temperature prediction and urban heat island analysis, including frameworks based on multisensor data and machine learning methods for prediction and assessment [23, 24]. However, there are still relatively few studies that combine satellite-derived thermal fields with daily meteorological forcing within a practical framework for next-day urban thermal prediction.

2.2 Urban heat risk assessment and vulnerability mapping

Another important direction is the transformation of thermal information into risk-oriented spatial analysis. Recent studies demonstrate that urban heat assessment should not be limited to temperature mapping but should also incorporate exposure, vulnerability, and population-related factors [25]. In this sense, GIS- and remote-sensing-based frameworks are often used to combine thermal indicators with environmental and demographic layers to cause urban heat vulnerability or risk maps [5, 6]. However, many of these studies are mostly descriptive or retrospective, while there are still fewer practical strategies that directly connect short-term thermal prediction with simplified urban heat risk mapping.

2.3 GPU acceleration in environmental deep learning

The third important direction is computational efficiency. Deep learning workflows in environmental applications usually incorporate large spatiotemporal datasets, repeated model training, and computationally demanding mapping procedures, which makes acceleration important in practice [19]. In urban heat applications, GPU-based computation has already been utilized to accelerate heat-exposure mapping, and some studies documented reductions in computation time of more than 99% for specific spatial estimations [21]. In a broader sense, fast deduction is also essential for near-real-time prediction and repeated evaluation [26]. However, GPU acceleration is still often treated only as an implementation detail in the background, and less often as an explicit part of an integrated framework for urban thermal prediction and risk-oriented mapping. In general, the literature suggests that deep learning-based heat prediction [27, 28], urban heat risk assessment [29], and GPU-accelerated computation [26] have mostly been examined separately.



(a) Study area: Sarajevo Canton.



(b) Terrain-based spatial context of Sarajevo.

Fig. 1: Spatial context of the study area. Panel (a) shows the geographic extent of Sarajevo used in this chapter, while panel (b) illustrates the surrounding topographic structure that contributes to spatial variability in urban thermal behavior.

3 Study Area, Data Sources, and Preprocessing


3.1 Study area: Sarajevo

The case study for urban heat analysis encloses Sarajevo and its surrounding neighborhoods,


which were picked because their thermal behavior is influenced by a dense urban structure, valley topography, heterogeneous land cover, and ongoing urban development. In Fig. 1(a), the studied area combines urban zones, residential districts with distinct construction densities, central transport corridors, and encircling peri-urban space. Sarajevo and its surrounding urban areas are suitable for studying heat-related phenomena, as summer heat may be further intensified by limited air circulation and local atmospheric conditions. This is also endorsed by the terrain context presented in Fig. 1(b), where the city is encountered within a valley system and influenced by the surrounding elevated terrain. Such a spatial configuration makes Sarajevo a useful example for studying the relationship between broader meteorological forcing and local urban thermal response.

3.2 Satellite-derived thermal data: MODIS LST

The main thermal target considered in this paper is land surface temperature obtained from the MODIS sensor. MODIS LST products are widely used in urban climate and environmental studies because they provide spatially explicit thermal observations with regular temporal coverage. In this study, MODIS-derived thermal maps were used to represent day-to-day urban thermal conditions over Sarajevo. Some representative examples of summer thermal fields extracted from the prepared dataset are shown in Fig. 2, where the spatial temperature variability over the study area is clearly visible across different dates. Each valid MODIS scene was transformed into a spatial temperature field over Sarajevo and stored in a standardized form suitable for predictive modeling. After cropping to the study area and removing invalid observations where necessary, all target fields were arranged on a common 32×32 grid. In this way, a compact but still spatially informative thermal representation was obtained, suitable for deep learning and able to preserve the main spatial differences important for next-day urban thermal prediction.



(a) MODIS LST field over Sarajevo on 2022-07-14.



(b) MODIS LST field over Sarajevo on 2017-07-31.

Fig. 2: Examples of MODIS-derived summer land surface temperature fields over Sarajevo. The panels show representative daily thermal maps on a 32×32 grid, highlighting the spatial variability of urban surface temperature patterns across different dates.

3.3 Meteorological forcing data: Open-Meteo

In addition to thermal imagery, the proposed predictive framework uses daily meteorological forcing variables from the Open-Meteo service. These variables were included because urban thermal




Fig. 3: Representative spatial distribution of the Sarajevo locations used to obtain daily meteorological forcing variables from the Open-Meteo service. Using multiple locations provides a more representative characterization of atmospheric conditions across the study area than relying on a single-point record.

behavior depends not only on land-surface conditions but also on surrounding atmospheric conditions. To obtain a more representative description of daily forcing over the city, meteorological data were collected from several locations in Sarajevo rather than using a single point. Their spatial distribution is illustrated in Fig. 3. The last set of daily variables possesses temperature-related indicators, humidity, precipitation, radiation, cloudiness, and wind characteristics. Hence, the predictive framework merges spatial thermal information from satellite measurements with time-aligned daily meteorological forcing. The usage of multiple sampling spots was vital, as Sarajevo has spatially distinct urban and topographic markers, and a single meteorological record could not properly capture the variability in causing conditions across the more comprehensive study area.




Fig. 4: Preprocessing pipeline from raw inputs to model-ready outputs.

3.4 Data preprocessing and dataset construction

A preprocessing stage was necessary in order to transform the heterogeneous raw data into a consistent dataset suitable for deep learning. Since the considered data sources differ in spatial resolution, temporal frequency, and format, all inputs had to be organized within a common spatial and temporal framework. The overall workflow is summarized in Fig. 4, which shows the main steps from raw data acquisition to the construction of samples ready for model training. For the thermal component, invalid MODIS observations were removed as needed, and the remaining scenes were converted into standardized daily raster fields over Sarajevo. For the meteorological component, daily forcing variables were synchronized with the dates of the available thermal observations. This temporal alignment step was important because each training sample had to connect predictor variables from a given temporal window with the target thermal field of the next day. After synchronization, the dataset was divided into input-output pairs for supervised learning. In the case of sequence-based models such as ConvLSTM, the data were further arranged into short temporal sequences so that the network could learn both spatial structure and temporal evolution. Additional preprocessing steps, such as normalization and scaling, were also applied to improve stability during training and numerical stability. As shown in Fig. 4, the final result of preprocessing was not only a collection of aligned rasters and tabular variables, but a structured dataset prepared for the considered predictive architectures. The resulting dataset construction pipeline provided

the basis for the experiments presented later in this chapter.

4 Proposed GPU-Accelerated Framework

4.1 Problem formulation

The objective of the proposed framework is to predict urban thermal conditions for the next day over the Sarajevo study area and to provide outputs that can later be used for simplified urban heat risk assessment. Let X_t present information at day t , including the recent thermal state of the city represented by spatial temperature fields and the corresponding meteorological forcing variables. The purpose of the forecast is to comprehend a mapping such as

$$f: X_t \mapsto \hat{Y}_{t+1},$$

where $\hat{Y}_{t+1}$ is the prediction for day $t + 1$. In other words, the model output is a raster-like temperature map representing the expected spatial distribution of thermal intensity across the study area for the following day. Depending on the considered architecture, the input may be given either as a channel-stacked spatial representation or as a short temporal sequence of previous observations. In all cases, the goal is not only to predict a single scalar temperature value but to reconstruct the complete spatial thermal pattern, which makes the problem inherently spatiotemporal. To achieve the same forecast goal, two model families were assessed: the first is a CNN-based model used as a baseline predictor of spatial thermal patterns, while the second is a ConvLSTM model made to capture both spatial and temporal dependencies. The CNN and ConvLSTM models were trained on the same multi-source dataset and evaluated by metrics such as mean absolute error (MAE), root mean square error (RMSE), and the coefficient of determination (R^2).

4.2 Prediction models

We analyze the CNN model in the most detail. We introduced this model in order to use a simple deep learning predictor for spatial thermal patterns. The main reason for using it is its compact, fully convolutional architecture, which can

process raster-like inputs while remaining simple enough for stable training and direct implementation. We assume that the input tensor is of size $B \times T \times C \times H \times W$, where B , T , and C are the batch size, sequence length, and the number of channels for one day, respectively. Also, inside the CNN model, the temporal and channel dimensions were merged before convolutional processing, so that the input is mapped into the shape of dimension $B \times (T \cdot C) \times H \times W$. Except for the mentioned, the network includes three 3×3 convolutional layers with padding 1, followed by ReLU activations, and a final 1×1 convolution, producing the predicted single-channel thermal map. For the hidden width equal to 64, the channel progression is

$$(T \cdot C) \rightarrow 64 \rightarrow 64 \rightarrow 32 \rightarrow 1.$$

Therefore, the CNN baseline follows the structure:

- Conv2D($T \cdot C$, 64, 3×3 , padding = 1) + ReLU,
- Conv2D(64, 64, 3×3 , padding = 1) + ReLU,
- Conv2D(64, 32, 3×3 , padding = 1) + ReLU,
- Conv2D(32, 1, 1×1).

Although the architecture lacks pooling layers, the spatial resolution is kept throughout the network, and the output is fixed to a 32×32 thermal matrix. However, temporal development is not explicitly modeled, except through the implicit information furnished by channel stacking. The main predictive model used in this chapter is a ConvLSTM architecture. Unlike the CNN baseline, this model keeps the temporal structure explicitly and processes the input in its original sequential form, that is, as a tensor of size $B \times T \times C \times H \times W$. At each time step, the current input and the previous hidden state are concatenated and passed through a 3×3 convolution with padding 1. In that way, the input, forget, output, and candidate gates of the ConvLSTM cell are generated. In the final experiments, the hidden dimension was set to 32. After the last time step, the final hidden state is passed to a convolutional prediction head. This head consists of one convolutional layer of size 3×3 with 32 output channels, followed by a ReLU activation function, and a final convolutional layer of size 1×1 that produces the predicted single-channel thermal map. Therefore, the prediction head has the following form:

- Conv2D(32, 32, 3×3 , padding = 1) + ReLU,
- Conv2D(32, 1, 1×1).




Fig. 5: CNN baseline architecture used for next-day urban thermal prediction.

Such a design allows the ConvLSTM model to capture both the spatial arrangement of urban heat patterns and their short-term temporal evolution. This is important in the case of persistent heatwave conditions. The structure of the CNN baseline is shown in Fig. 5. The model first merges the temporal and channel dimensions, then applies a fully convolutional mapping to obtain the predicted thermal field for the next day. The ConvLSTM architecture used in the final experiments is shown in Fig. 6. In contrast to the CNN baseline, it preserves the temporal structure and processes the input sequence recurrently before generating the final thermal map by a convolutional prediction head.

4.3 GPU acceleration strategy

Because the proposed framework involves repeated training and evaluation of multiple deep learning architectures on multi-source spatiotemporal data, computational efficiency becomes an important practical issue. For that reason, all experiments in this chapter are designed around GPU-accelerated execution. The experiments were carried out on a workstation equipped




Fig. 6: ConvLSTM architecture used in the final experiments.

with an AMD64 processor, 16 GB of RAM, and Windows 11. Training was performed on CUDA-enabled hardware, and GPU acceleration was used in order to reduce computation time and allow repeated model evaluation under different dataset configurations. The implementation relies on CUDA-enabled tensor computation, which allows convolutions and recurrent updates to be executed in parallel. This substantially reduces the training time compared with a purely CPU-based workflow. In addition, mixed-precision training is employed where appropriate in order to further improve computational efficiency and reduce memory usage. This is particularly useful when working with image-like spatial data and sequence-based models, since both can lead to significant memory demands during training. From a methodological perspective, GPU acceleration is not treated only as a technical implementation detail, but as an integral part of the proposed practical framework. Without accelerated training, repeated experimentation with CNN and ConvLSTM models would be significantly more difficult, especially when multiple input configurations and loss formulations must be tested. The

use of GPU resources therefore directly supports the experimental design of the chapter and helps make the proposed urban thermal prediction workflow computationally viable in practice.

5 Urban Heat Risk Assessment Layer

Although the primary focus of this chapter is next-day urban thermal prediction, the obtained predictive maps can also be used as the basis for a simplified urban heat risk assessment framework. In climate-risk literature, risk is commonly interpreted through the interaction of hazard, exposure, and vulnerability, and this perspective has also been adopted in urban heat studies [30, 31]. The main idea here is therefore to extend temperature prediction toward a more decision-oriented spatial interpretation by combining the predicted thermal field with additional exposure and vulnerability information. In this chapter, this risk-oriented extension is intentionally kept simple and transparent. Its purpose is not to provide a fully operational urban risk system, but rather to illustrate how GPU-accelerated thermal prediction can support practical urban heat analysis and spatial screening of potentially critical urban zones [32].

5.1 Hazard layer

The hazard layer represents the spatial intensity of predicted thermal stress over the study area. In the proposed framework, it is derived directly from the predicted next-day thermal field produced by the deep learning models. Such an interpretation is consistent with earlier urban heat risk studies in which thermal indicators, including land surface temperature or related heat-stress measures, are used as the hazard component of the risk framework [30, 32]. Locations with higher predicted temperature values are interpreted as having higher thermal hazard, while locations with lower values correspond to lower hazard. For integration with the remaining components of the framework, the predicted field can be normalized to a common interval such as $[0, 1]$, which is a standard step when combining heterogeneous indicators with different units and scales [32, 33]. In this way, the hazard layer serves as the predictive component of the urban heat risk assessment,

indicating where thermally critical conditions are expected to occur on the following day.

5.2 Exposure layer

The exposure layer is introduced to represent the spatial distribution of people or urban elements that may be affected by elevated thermal conditions. Its role is important because high thermal intensity alone does not necessarily indicate a high practical risk, unless the affected area also includes a considerable human presence or urban activity. In urban heat risk mapping, exposure is often represented by population density, pedestrian presence, residential concentration, or other indicators of human activity [30, 32]. In a simplified setting, the exposure layer can therefore be constructed from gridded demographic data or from proxy indicators such as residential density or built-up intensity. After alignment with the spatial grid of the hazard layer, the obtained surface is normalized so that higher values indicate stronger exposure [33]. In this way, it becomes possible to distinguish between areas with low exposure to heat and those where heat may affect a larger number of people.

5.3 Vulnerability layer

The vulnerability layer describes the differential sensitivity of different parts of the city to thermal stress. In urban heat research, vulnerability is commonly associated with environmental, demographic, social, and urban-form characteristics that increase sensitivity or reduce coping capacity under hot conditions [33–35]. In real life, these elements may include limited vegetation, dense built-up structure, impervious surfaces, low ventilation potential, disadvantaged socioeconomic conditions, or the presence of particularly sensitive population groups. Several recent studies have shown that vulnerability-oriented heat assessments benefit from integrating both social indicators and spatial characteristics of the urban fabric [34, 35]. In the present chapter, vulnerability is represented in a simplified manner through spatial indicators that align with the study grid. As in the case of hazard and exposure, the resulting layer is normalized before integration [33]. Higher values indicate locations where predicted heat is expected to have stronger adverse implications due to less favorable local conditions.

5.4 Risk index generation

After the hazard, exposure, and vulnerability layers have been organized, they can be united into one simplified urban heat risk index. This follows the general idea of climate- and urban heat-risk assessment, where the most critical areas are those in which a strong thermal hazard co-occurs with high exposure and increased vulnerability [32, 33]. One simple and transparent formulation is

$$R(i, j) = H(i, j) \cdot E(i, j) \cdot V(i, j),$$

where $H(i, j)$, $E(i, j)$, and $V(i, j)$ denote the normalized hazard, exposure, and vulnerability values at spatial position (i, j) , respectively. Multiplicative and similar composite formulations based on normalized indicators are often used in heat-vulnerability and heat-risk studies because they emphasize the joint effects of multiple factors, rather than relying solely on temperature [32, 33]. The final risk surface can then be presented as a raster map or divided into ordinal categories such as low, moderate, high, and very high risk. Although this formulation is simplified, it is sufficient to show how next-day thermal forecasting can be extended towards interpretable urban heat risk assessment products [30, 33].

6 Experimental Results

The experimental part of this article was designed as a practical case study for next-day urban thermal prediction over Sarajevo. The goal was to predict the next daily land surface temperature (LST) map from a short temporal sequence of previous observations and meteorological forcing variables. All experiments were performed on CUDA-enabled hardware, and GPU acceleration was used during training in order to reduce training time and allow repeated model evaluation under different dataset configurations. The target output in all experiments was a normalized MODIS LST map of spatial size 32×32 . Hence, each target map contains 1024 pixel values. The input tensor for one sample had the general form

$$X \in \mathbb{R}^{T \times C \times H \times W},$$

where T denotes the sequence length, C denotes the number of channels per day, and $H = W = 32$.

In the final experiments, the sequence length was fixed to $T = 3$. Two main dataset configurations were considered. The first configuration combined daily MODIS LST maps with daily meteorological forcing extracted from a single Open-Meteo location over Sarajevo. The second configuration extended both the temporal coverage and the number of meteorological inputs by using four representative Sarajevo locations together with a longer time interval. The final configuration of the best-performing ConvLSTM experiment is summarized in Table 1. It includes the main characteristics of the input representation, the temporal setup, the training configuration, and the evaluation protocol used in the final stage of the study. All experiments were carried out on a workstation equipped with an AMD64 processor, 16 GB of RAM, and Windows 11.

Table 1: Final experimental setup used for the best-performing ConvLSTM model.

Setting	Value
Input data	MODIS LST Open-Meteo daily variables
Study period	2010–2025
Meteorological locations	4
Input channels per day	41
Sequence length	3
Input shape	$3 \times 41 \times 32 \times 32$
Output shape	$1 \times 32 \times 32$
Valid samples	1871
Train/validation split	80% / 20%
Model	ConvLSTM
Hidden dimension	32
Optimizer	Adam
Learning rate	0.001
Batch size	16
Epochs	20
Loss	$0.7 L_1 + 0.3 L_2$
Training mode	GPU + mixed precision
Metrics	MAE, RMSE, R^2

6.1 Dataset Construction

The first dataset configuration was obtained by combining MODIS MOD11A1 daily daytime LST data with Open-Meteo daily meteorological variables for Sarajevo. In the case of one-location forcing, the following meteorological variables were used:

- temperature_2m_mean,
- temperature_2m_max,
- temperature_2m_min,
- dew_point_2m_mean,
- precipitation_sum,
- wind_speed_10m_mean,
- shortwave_radiation_sum.

These variables were combined with one MODIS LST channel, so that the total number of channels per day was equal to 8. For the temporal sequence length equal to 3, one input sample had the form

$$3 \times 8 \times 32 \times 32.$$

After preprocessing and removal of invalid cases, this baseline configuration contained 667 valid samples. In the second configuration, the dataset was extended in several ways. First, the MODIS period was expanded to 2010–2025, while the initial configuration covered the period 2019–2024. Second, meteorological forcing was not collected from only one point, but from four Sarajevo locations: Sarajevo Center, Ilidža, Dobrinja, and Vogošća. Third, the number of meteorological variables was increased to include:

- temperature_2m_mean,
- temperature_2m_max,
- temperature_2m_min,
- dew_point_2m_mean,
- precipitation_sum,
- shortwave_radiation_sum,
- relative_humidity_2m_mean,
- cloud_cover_mean,
- wind_direction_10m_dominant,
- wind_speed_10m_mean.

Since these variables were taken for four locations, the meteorological forcing contributed 40 channels per day. Together with the MODIS LST channel, the total number of channels per day was 41. Therefore, one input sample had the form

$$3 \times 41 \times 32 \times 32.$$

After alignment in time and removal of invalid cases, the final extended dataset contained 1871 valid samples.

6.2 Models and Training Procedure

Two deep learning architectures were considered during the experiments: a CNN baseline and a

ConvLSTM model. The CNN was used as a simpler reference model, while the ConvLSTM was expected to better capture spatiotemporal dependencies. In the final stage of the experiments, the ConvLSTM model clearly showed better behaviour and was therefore selected as the main model for detailed analysis. The training and validation sets were obtained using an 80%/20% split. Thus, for the extended dataset configuration, the model was trained on 1538 samples and validated on 333 samples. In all experiments, training was performed on GPU. Mixed precision was also used during training in order to improve computational efficiency. At first, a standard mean squared error loss was used. However, the predicted maps were often too smooth, even when the quantitative metrics were acceptable. For this reason, a hybrid loss was introduced:

$$\mathcal{L} = \alpha \mathcal{L}_1 + (1 - \alpha) \mathcal{L}_2,$$

where $\mathcal{L}_1$ denotes the mean absolute error loss, $\mathcal{L}_2$ denotes the mean squared error loss, and $\alpha = 0.7$ in our experiments. This modification led to both better numerical results and visually more realistic thermal maps.

6.3 Evaluation Metrics

The models were evaluated using three standard regression metrics:

- mean absolute error (MAE),
- root mean squared error (RMSE),
- coefficient of determination (R^2).

These metrics were computed over the predicted and reference LST maps.

6.4 Quantitative Results

The most important quantitative results are summarized in Table 2. The first strong result was obtained on the MODIS + Open-Meteo single-location dataset 1 with 667 samples. Using the ConvLSTM model with the hybrid loss, the method achieved the following values: **MAE** = 0.2543, **RMSE** = 0.3349, and **R^2** = 0.8606. After extending the temporal range and replacing the single-location forcing with a multi-location meteorological representation, the results improved further. On the final dataset 2 with 1871 samples,

the ConvLSTM model achieved the following values: **MAE** = 0.2293, **RMSE** = 0.3089, and **R^2** = 0.8877. These results show that both a stronger dataset and a richer meteorological representation improve next-day urban thermal prediction.

Table 2: Main quantitative results obtained with the ConvLSTM model.

Dataset configuration	Num. of samples	MAE	RMSE	R^2
Dataset 1	667	0.2543	0.3349	0.8606
Dataset 2	1871	0.2293	0.3089	0.8877

6.5 Qualitative Results

Besides the global metrics, qualitative inspection of predicted maps was also performed. In general, the ConvLSTM model was able to recover the main large-scale thermal structure of the target maps. The warm regions and the main spatial gradients were usually well reconstructed, especially in representative medium-complexity cases. However, the predictions were still smoother than the reference MODIS maps. This means that the model was more successful in capturing coarse spatial patterns than very fine local anomalies. Such behavior is expected, since the target maps contain sharp local irregularities, while the model tries to learn stable spatiotemporal structure from limited and partially noisy observations. Figure 7 presents three representative qualitative examples. The first row contains the reference thermal maps, while the second row shows the corresponding predictions produced by the ConvLSTM model. In the first two examples, the model follows the main spatial structure of the reference maps reasonably well and produces relatively small prediction errors. The corresponding errors for the predicted maps are as follows: first predicted map (MAE = 0.1548, RMSE = 0.2387) and second predicted map (MAE = 0.1319, RMSE = 0.1868). In the third example, however, the model fails to reproduce a more difficult and strongly localized thermal pattern, which results in substantially larger error values (MAE = 1.1496, RMSE = 1.1693). Overall, the qualitative analysis confirms that the model is generally able to capture the dominant thermal structure of the target maps,

although the predicted maps remain smoother than the corresponding reference MODIS maps. To better understand the global behavior of the model over the full evaluation set, the per-sample MAE values were analyzed in two complementary ways. Figure 8(a) shows the distribution of per-sample MAE values. It can be seen that most samples are concentrated in the low-to-moderate error range, while only a relatively small number of cases produce large errors. This indicates that the overall performance of the model is generally stable, although a limited number of difficult outliers still remain. Figure 8(b) presents the MAE value for each evaluated sample. The plot shows that, for most samples, the error remains relatively low, while occasional peaks correspond to more challenging thermal patterns that are harder for the model to reproduce accurately.

6.6 Discussion

The experiments clearly indicate that dataset quality plays a major role in the final predictive performance. The initial experiments, which relied on a smaller and less informative dataset, produced smoother predictions and less stable quantitative results. After increasing the number of samples, extending the temporal coverage, and enriching the meteorological forcing with data from multiple Sarajevo locations, the overall performance improved substantially. The best results were achieved by the final ConvLSTM model trained with the hybrid loss function. Its R^2 score of 0.8877 indicates a strong agreement between the predicted and reference thermal maps. At the same time, the qualitative analysis shows that some limitations still remain. In particular, the model tends to smooth out very fine local thermal details, and a limited number of challenging cases still produce relatively large errors. Nevertheless, the obtained results are strong enough to confirm the usefulness of the proposed practical framework. They show that a GPU-accelerated ConvLSTM model, combined with satellite-derived thermal observations and enriched daily meteorological forcing, can provide meaningful next-day urban thermal predictions for a real urban environment.




Fig. 7: Qualitative comparison between reference and predicted thermal maps for three representative examples.




Fig. 8: (a) Distribution of per-sample MAE values. (b) Per-sample MAE across all evaluated samples.

7 Conclusion

This chapter presented a practical GPU-accelerated deep learning framework for next-day urban thermal prediction and simplified urban heat risk assessment. The proposed approach combined MODIS land surface temperature maps with daily meteorological forcing variables to

model short-term thermal dynamics over Sarajevo. Special attention was given to multi-source dataset construction, spatiotemporal deep learning models, and the use of GPU acceleration to enable efficient repeated training and evaluation. The results showed that dataset design strongly affects predictive performance. While early experiments based on smaller datasets produced

smoother and less stable predictions, substantial improvements were achieved after extending the temporal range, increasing the number of valid samples, and introducing multi-location meteorological forcing. The best performance was obtained with the ConvLSTM model trained using a hybrid loss function, reaching an MAE of 0.2293, an RMSE of 0.3089, and an R^2 score of 0.8877. These results indicate strong agreement between predicted and reference thermal maps, although some fine local details remain difficult to reconstruct. Overall, the findings confirm that GPU-accelerated deep learning provides a practical basis for urban thermal prediction in real environments. The proposed framework is flexible enough to support further methodological improvements and application-oriented extensions related to urban climate adaptation, early warning, and heat risk assessment.

References

- [1] Cheval, S., Amihăsesei, V.-A., Chitu, Z., Dumitrescu, A., Falcescu, V., Iraşoc, A., Micu, D.M., Mihuleţ, E., Ontel, I., Paraschiv, M.-G., Tudose, N.C.: A systematic review of urban heat island and heat waves research (1991–2022). *Climate Risk Management* **44**, 100603 (2024) <https://doi.org/10.1016/j.crm.2024.100603>
- [2] Cuervo-Vilches, T., Díaz, J., López-Bueno, J.A., Luna, M.Y., Navas, M.A., Mirón, I.J., Linares, C.: Impact of urban heat islands on morbidity and mortality in heat waves: Observational time series analysis of Spain's five cities. *Science of The Total Environment* **890**, 164412 (2023) <https://doi.org/10.1016/j.scitotenv.2023.164412>
- [3] Huang, W.T.K., Masselot, P., Bou-Zeid, E., Fatichi, S., Paschalis, A., Sun, T., Gasparini, A., Manoli, G.: Economic valuation of temperature-related mortality attributed to urban heat islands in European cities. *Nature Communications* **14**(1), 1–12 (2023) <https://doi.org/10.1038/s41467-023-43135-z>
- [4] Hsu, A., Sheriff, G., Chakraborty, T., Manya, D.: Disproportionate exposure to urban heat island intensity across major US cities. *Nature Communications* **12**(1), 1–11 (2021) <https://doi.org/10.1038/s41467-021-22799-5>
- [5] Pan, X., Mavrokapnidis, D., Ly, H.T., Mohammadi, N., Taylor, J.E.: Assessing and forecasting collective urban heat exposure with smart city digital twins. *Scientific Reports* **14**(1), 1–14 (2024) <https://doi.org/10.1038/s41598-024-59228-8>
- [6] D'Ambrosio, V., Martino, F.D., Miraglia, V.: A GIS-based framework to assess heatwave vulnerability and impact scenarios in urban systems. *Scientific Reports* **13**(1) (2023) <https://doi.org/10.1038/s41598-023-39820-0>
- [7] Ye, T., Xu, R., Huang, W., Yang, Z., Yu, P., Yu, W., Liu, Y., Wu, Y., Wen, B., Zhang, Y., Hart, J.E., Nieuwenhuijsen, M., Abramson, M.J., Guo, Y., Li, S.: Billions of people exposed to increasing heat but decreasing greenness from 2000 to 2022. *The Innovation* **6**(5), 100870 (2025) <https://doi.org/10.1016/j.xinn.2025.100870>
- [8] Tatem, A.J.: WorldPop, open data for spatial demography. *Scientific Data* **4**, 1–4 (2017) <https://doi.org/10.1038/sdata.2017.4>
- [9] Wan, Z.: New refinements and validation of the MODIS land-surface temperature/emissivity products. *Remote Sensing of Environment* **112**(1), 59–74 (2008) <https://doi.org/10.1016/j.rse.2006.06.026>
- [10] Duan, S.-B., Li, Z.-L., Li, H., Göttsche, F.-M., Wu, H., Zhao, W., Leng, P., Zhang, X., Coll, C.: Validation of collection 6 MODIS land surface temperature product using in situ measurements. *Remote Sensing of Environment* **225**, 16–29 (2019) <https://doi.org/10.1016/j.rse.2019.02.020>
- [11] Lloyd, C.T., Sorichetta, A., Tatem, A.J.: High resolution global gridded data for use in population studies. *Scientific Data* **4**(1), 170001 (2017) <https://doi.org/10.1038/sdata.2017.1>
- [12] Barriopedro, D., García-Herrera, R., Ordóñez, C., Miralles, D.G., Salcedo-Sanz,

- S.: Heat waves: Physical understanding and scientific challenges. *Reviews of Geophysics* **61**(2), 2022–000780 (2023) <https://doi.org/10.1029/2022RG000780>
- [13] Xu, C., Wei, R., Tong, H.: A systematic review of methodological advances in urban heatwave risk assessment: Integrating multi-source data and hybrid weighting methods. *Sustainability* **17**(8) (2025) <https://doi.org/10.3390/su17083747>
- [14] Boudreault, J., Lamothe, F., Campagna, C., Chebana, F.: Machine learning for modelling the health impacts of extreme heat: A comprehensive literature review. *Environment International* **206**, 109965 (2025) <https://doi.org/10.1016/j.envint.2025.109965>
- [15] Zhang, H., Liu, Y., Zhang, C., Li, N.: Machine learning methods for weather forecasting: A survey. *Atmosphere* **16**(1) (2025) <https://doi.org/10.3390/atmos16010082>
- [16] Tuba, E., Tuba, I., Hrosik, R.C., Alihodzic, A., Tuba, M.: Image classification by optimized convolution neural networks. In: Rathore, V.S., Sharma, S.C., Tavares, J.M.R.S., Moreira, C., Surendiran, B. (eds.) *Rising Threats in Expert Applications and Solutions*, pp. 447–454. Springer, Singapore (2022)
- [17] Shekhar, S., Jiang, Z., Ali, R.Y., Eftelioglu, E., Tang, X., Gunturi, V.M.V., Zhou, X.: Spatiotemporal data mining: A computational perspective. *ISPRS International Journal of Geo-Information* **4**(4), 2306–2338 (2015) <https://doi.org/10.3390/ijgi4042306>
- [18] Niño Medina, J.S., Suarez Barón, M.J., Reyes Suarez, J.A.: Application of deep learning for the analysis of the spatiotemporal prediction of monthly total precipitation in the boyacá department, colombia. *Hydrology* **11**(8) (2024) <https://doi.org/10.3390/hydrology11080127>
- [19] Pandey, M., Fernandez, M., Gentile, F., Isayev, O., Tropsha, A., Stern, A.C., Cherkasov, A.: The transformational role of gpu computing and deep learning in drug discovery. *Nature Machine Intelligence* **4**(3), 211–221 (2022) <https://doi.org/10.1038/s42256-022-00463-x>
- [20] Mwitta, C., Rains, G.C., Prostko, E.: Evaluation of inference performance of deep learning models for real-time weed detection in an embedded computer. *Sensors* **24**(2) (2024) <https://doi.org/10.3390/s24020514>
- [21] Li, X., Wang, G.: Gpu parallel computing for mapping urban outdoor heat exposure. *Theoretical and Applied Climatology* **145**(3), 1101–1111 (2021) <https://doi.org/10.1007/s00704-021-03692-z>
- [22] Shafiq, F., Zafar, A., Khan, M.U.G., Iqbal, S., Albeshier, A.S., Asghar, M.N.: Extreme heat prediction through deep learning and explainable ai. *PLOS One* **20**(3) (2025) <https://doi.org/10.1371/journal.pone.0316367>
- [23] Tang, N., Farhan, M., Mohammad, P., Abdullah-Al-Wadud, M., Hussain, S., Hamza, U., Zulqarnain, R.M., Rebouh, N.Y.: Enhancing urban heat island analysis through multisensor data fusion and gru-based deep learning approaches for climate modeling. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **18**, 9279–9296 (2025) <https://doi.org/10.1109/JSTARS.2025.3554529>
- [24] Wang, L.: Urban heat island monitoring and intelligent prediction method based on big data analysis. *Sustainable Energy Research* **12**(1) (2025) <https://doi.org/10.1186/s40807-025-00206-7>
- [25] Li, F., Yigitcanlar, T., Nepal, M., Thanh, K.N., Dur, F.: A novel urban heat vulnerability analysis: Integrating machine learning and remote sensing for enhanced insights. *Remote Sensing* **16**(16) (2024) <https://doi.org/10.3390/rs16163032>
- [26] Kalfarisi, R., Chew, A.W.Z., Cai, J., Xue, M., Pok, J., Wu, Z.Y.: Predictive modeling framework accelerated by gpu computing for smart water grid data-driven analysis in near real-time. *Advances in Engineering Software* **173**, 103287 (2022) <https://doi.org/10.1016/>

- [27] Ge, S., Zhan, W., Li, J., Li, L., Dong, P., Li, X., Wang, C., Wang, C., Gao, Y.: Prediction of next-day 1-km canopy urban heat island by integrating a multi-block convolutional neural network with satellite- and ground-based observations. *Sustainable Cities and Society* **131**, 106701 (2025) <https://doi.org/10.1016/j.scs.2025.106701>
- [28] Lyu, F., Wang, S., Han, S.Y., Catlett, C., Wang, S.: An integrated cyberGIS and machine learning framework for fine-scale prediction of urban heat island using satellite remote sensing and urban sensor network data. *Urban Informatics* **1**, 6 (2022) <https://doi.org/10.1007/s44212-022-00002-4>
- [29] Emmanuel, R., Jalal, M., Ogunfuyi, S., Maharoof, N., Zala, M., Perera, N., Ratnayake, R.: Urban heat risk: Protocols for mapping and implications for Colombo, Sri Lanka. *Atmosphere* **14**(2) (2023) <https://doi.org/10.3390/atmos14020343>
- [30] Colaninno, N., Basu, R., Hosseini, M., Alhassan, A., Liu, L., Sevtsuk, A.: A sidewalk-level urban heat risk assessment framework using pedestrian mobility and urban microclimate modeling. *Environment and Planning B: Urban Analytics and City Science* **52**(5), 1071–1090 (2025) <https://doi.org/10.1177/23998083241280746> <https://doi.org/10.1177/23998083241280746>
- [31] Morabito M, G.B.G.G.T.P.D.S.V.O.S.G.G. Crisci A: Urban-hazard risk analysis: mapping of heat-related risks in the elderly in major Italian cities. *PLoS One* **18**(5) (2015) <https://doi.org/10.1371/journal.pone.0127277>
- [32] Zhang, W., Zheng, C., Chen, F.: Mapping heat-related health risks of elderly citizens in mountainous area: A case study of Chongqing, China. *Science of The Total Environment* **663**, 852–866 (2019) <https://doi.org/10.1016/j.scitotenv.2019.01.240>
- [33] Wu, C., Shui, W., Huang, Z., Wang, C., Wu, Y., Wu, Y., Xue, C., Huang, Y., Zhang, Y., Zheng, D.: Urban heat vulnerability: A dynamic assessment using multi-source data in coastal metropolis of southeast China. *Frontiers in Public Health* **Volume 10 - 2022** (2022) <https://doi.org/10.3389/fpubh.2022.989963>
- [34] Voelkel, J., Hellman, D., Sakuma, R., Shandas, V.: Assessing vulnerability to urban heat: A study of disproportionate heat exposure and access to refuge by socio-demographic status in Portland, Oregon. *International Journal of Environmental Research and Public Health* **15**(4) (2018) <https://doi.org/10.3390/ijerph15040640>
- [35] Técher, M., Ait Haddou, H., Aguejdad, R.: Urban heat island's vulnerability assessment by integrating urban planning policies: A case study of Montpellier Méditerranée Métropolitaine area, France. *Sustainability* **15**(3) (2023) <https://doi.org/10.3390/su15031820>